

A simple algorithm for random colouring $G(n, d/n)$ using $(2 + \epsilon)d$ colours. *

Charilaos Efthymiou

University of Warwick, Mathematics and Computer Science, Coventry CV4 7AL, UK
c.efthymiou@warwick.ac.uk

July 6, 2011

Abstract

Approximate random k -colouring of a graph $G = (V, E)$ is a very well studied problem in computer science and statistical physics. It amounts to constructing a k -colouring of G which is distributed close to *Gibbs distribution*, i.e. the uniform distribution over all the k -colourings of G . Here, we deal with the problem when the underlying graph is an instance of Erdős-Rényi random graph $G(n, p)$, where $p = d/n$ and d is fixed.

We propose a novel efficient algorithm for approximate random k -colouring with the following properties: given an instance of $G(n, d/n)$ and for any $k \geq (2 + \epsilon)d$, it returns a k -colouring distributed within total variation distance $n^{-\Omega(1)}$ from the Gibbs distribution, with probability $1 - n^{-\Omega(1)}$.

What we propose is neither a MCMC algorithm nor some algorithm inspired by the message passing heuristics that were introduced by statistical physicists. Our algorithm is of *combinatorial* nature. It is based on a rather simple recursion which reduces the random k -colouring of $G(n, d/n)$ to random k -colouring simpler subgraphs first.

The lower bound on the number of colours for our algorithm to run in polynomial time is *dramatically* smaller than the corresponding bounds we have for any previous algorithm.

Key words: Random colouring, sparse random graph, efficient algorithm.

*Supported by EPSRC grant EP/G039070/2 and DIMAP.

1 Introduction

Approximate random k -colouring of a graph $G = (V, E)$ is a very well studied problem in computer science and statistical physics. It amounts to constructing a k -colouring of G which is distributed close to *Gibbs distribution*, i.e. the uniform distribution over all the k -colourings of G . Here, we deal with the specific algorithmic problem when the underlying graph is an instance of Erdős-Rényi random graph $G(n, p)$, where $p = d/n$ and d is fixed.

The most powerful and most popular algorithms for this kind of problems are based on the Markov Chain Monte Carlo (MCMC) method. There the main technical challenge is to establish that the underlying Markov chain mixes in polynomial time (see [6]). The work in [8] (which improved [3]) shows that the well known Markov chain *Glauber dynamics* for k -colourings has polynomial time mixing for typical instances of $G(n, d/n)$ as long as the number of colours k is larger than $k(d)$, a number which depends only on the expected degree of $G(n, d/n)$, d .

Notably, both [8, 3] overcame the “maximum degree obstacle” from which most techniques for analysing the mixing time of the Glauber dynamics suffer, i.e. they are stated in terms of the maximum degree of the graph. This makes them insufficient for $G(n, d/n)$ where the maximum degree grows as $\Theta((\log n)/(\log \log n))$ with probability $1 - o(1)$.

Recently physicists proposed some heuristics for computing deterministically marginals of Gibbs distribution. These heuristics are based on the message passing algorithm Belief Propagation (see [7]) and *ideally* they could be used for random colouring $G(n, d/n)$. There, the main challenge is to show that the computation of the marginals is accurate. In turn this amounts to establishing certain *spatial* correlation decay conditions for the Gibbs distribution. We should remark that the heuristics proposed by statistical physicists, namely *Belief Propagation guided decimation* and *Survey Propagation guided decimation*, were put forward on the basis of very insightful but highly non-rigorous statistical mechanics considerations (see [2]). The work in [4] presents an efficient algorithm which is a variation of Belief propagation and returns an approximate random k -colouring of a typical $G(n, d/n)$ as long as $k > k'(d)$, where $k'(d)$ is a number which depends only the expected degree (i.e. $k'(d) = d^{14}$).

In this work we propose a novel algorithm for approximate random k -colouring $G(n, d/n)$ which not only overcomes the “maximum degree obstacle” but somehow *optimizes* the dependence of the minimum number of colours from the expected degree d . The lower bound on the number of colours for our algorithm to run in polynomial time is *dramatically* smaller than the corresponding bounds we have for any previous algorithm on the problem. The algorithm does not fall into any of the previous two categories, i.e., it is neither MCMC nor based on Gibbs marginals computation.

We are based on the following humble observation: Let $\{v, u\}$ be an edge of $G_{n, d/n}$. A random colouring of $G_{n, d/n}$ can be seen as a random colouring of $G_{n, d/n} \setminus \{v, u\}$ with the additional property that v and u are assigned different colours. Assume that we have a polynomial time algorithm, call it STEP, such that given any graph G and two non-adjacent vertices v, u it transforms a random colouring of G to a random colouring which has the *extra* property that v and u take different colours. In that case, the initial problem can be reduced to taking a random k -colouring of $G_{n, d/n} \setminus \{v, u\}$ and then use STEP. The reasonable question, then, would be how can someone colour $G_{n, d/n} \setminus \{v, u\}$ randomly. We can set up a recursion by applying the previous reduction for $G_{n, d/n} \setminus \{v, u\}$ and so on. Note that as the recursion proceeds the structure of the graph that is considered gets simpler and simpler. This is due to the edge deletions. Clearly, after a certain number of recursive calls the graph becomes so simple that it can be k -coloured randomly in polynomial time by some known algorithm.

A great deal of this work illustrates the implementation of STEP in the special case where the input graph G is a typical instance of $G_{n, d/n}$ or any of its subgraphs that are considered in the recursion. STEP will be an approximation algorithm, i.e. given a random colouring of G in the input the distribution of the output will be an approximation of the desired one. Consequently, at the end we get an approximate random k -colouring of $G(n, d/n)$.

We use total variation distance as a measure of distance between distributions.

Definition 1 For the distributions ν_a, ν_b on $[k]^V$, let $\|\nu_a - \nu_b\|$ denote their total variation distance, i.e.

$$\|\nu_a - \nu_b\| = \max_{\Omega' \subseteq [k]^V} |\nu_a(\Omega') - \nu_b(\Omega')|.$$

For $\Lambda \subseteq V$ let $\|\nu_a - \nu_b\|_\Lambda$ denote the total variation distance between the projections of ν_a and ν_b on $[k]^\Lambda$.

STEP will have the following *general property*: Consider in the input a random k -colouring of some graph G and v, u , two non-adjacent vertices of G . The accuracy of the outcome depends on certain *spatial mixing* properties of the Gibbs distribution of the colourings of G . In particular, for a random k -colouring of G it suffices that there is a sufficiently large $b > 0$ such that

$$\left| \Pr[u \text{ is coloured } c | v \text{ is coloured } q] - \frac{1}{k} \right| \leq \exp(-b \cdot \text{dist}(v, u)) \quad \forall c, q \in [k]. \quad (1)$$

Moreover, assuming that (1) holds, then the distribution of the output of STEP is within total variation distance from the ideal distribution a quantity which is proportional to the r.h.s. of (1). Consequently, when we consider the previous recursive random colouring algorithm (that uses STEP), we note that it is desirable to delete edges that belong to long cycles in each recursive call.

We show that for a typical $G(n, d/n)$ and for $k \geq (2 + \epsilon)d$, where $\epsilon > 0$ is fixed, we get a relation as in (1) for the random k -colourings of any graph in the recurrence. Moreover, if we are careful enough on how do we delete the edges in the recurrence, the outcome of the random colouring algorithm is very close to Gibbs distribution. In particular, we show the following theorem.

Theorem 1 Let μ be the uniform distribution over the k -colourings of $G_{n,d/n}$ and let μ' be the distribution of the colouring that is returned by our random colouring algorithm. Taking $k \geq (2 + \epsilon)d$, for fixed $\epsilon > 0$, then with probability at least $1 - n^{-\frac{\epsilon}{90 \log d}}$ it holds that

$$\|\mu - \mu'\| = O\left(n^{-\frac{\epsilon}{90 \log d}}\right),$$

for any fixed $d > d_0(\epsilon)$.

Additionally, we provide guarantees on the time complexity of the algorithm.

Theorem 2 With probability at least $1 - n^{-2/3}$, it holds that the time complexity of the random colouring algorithm is $O(n^2)$.

Detailed proofs of Theorem 1 and Theorem 2 appear in the appendix, Section A.

Notation We denote with small letters of the greek alphabet the colourings of a graph G , e.g. σ, η, τ , while we use capital letters for the random variables which take values over the colourings e.g. X, Y, Z . We denote with σ_v the colour assignment of the vertex v under the colouring σ . Similarly, the random variable $X(v)$ is equal to the colour assignment that X specifies for the vertex v . Finally, for an integer $k > 0$ let $[k] = \{1, \dots, k\}$.

2 Basic Description

In this section we provide a more detailed description of our approximate colouring algorithm. We assume that the input graph is an instance of $G(n, d/n)$ and k is the numbers of colours.

Set up. Consider a sequence of graphs G_0, \dots, G_r such that every G_i is a subgraph of $G_{n,d/n}$. Moreover, G_r is identical to $G_{n,d/n}$, while G_i is derived by deleting some edge of G_{i+1} .

So as to get the graph G_i from G_{i+1} the only rule we follow is that we delete, arbitrarily, an edge that belongs to a sufficiently large cycle, i.e of length at least $(\log n)/(9 \log d)$. G_0 is the graph that comes up when there no are other such edges to delete. Note only that G_i , as a subgraph of $G(n, d/n)$, is somehow *random*.

Colouring. With probability $1 - n^{-\Omega(1)}$, the sequence of subgraphs has the property that G_0 is simple enough and we can k -colour it randomly in polynomial time by using some known algorithm. In that case the algorithm takes a random colouring of G_0 . Then, for $i = 0$ to $r - 1$ it does the following: it takes the random colouring of G_i , it does a simple, i.e. polynomial time, processing of this colouring and gets a random colouring of G_{i+1} . The algorithm continues until G_r .

Let G and G' be two consecutive terms in the sequence of graphs, above. Assume that G is derived by deleting the edge $\{v, u\}$ from G' . The critical question is the following one: Given X , a random k -colouring of G , how can someone use it to get efficiently X' , a random k -colouring of G' . A moment's reflection makes it clear that if X has the additional property that $X(v) \neq X(u)$, then X is distributed u.a.r. among the k -colourings of G' . In this case we can simply set $X' = X$. Unfortunately, this cannot always be the case and the random colouring algorithm we propose somehow deals with situations as the one where $X(v) = X(u)$.

Definition 2 (Good & Bad colourings) Let σ be a proper k -colouring of G . We call σ a bad colouring of G if $\sigma_v = \sigma_u$. Otherwise, we call σ a good colouring of G .

It turns out that the basic algorithmic challenge here is captured in the following problem.

Problem 1 Given a bad random colouring of G , turn it to a good random colouring, in polynomial time.

Let us give an intuitive description of our algorithm for the above problem. First remark the following: Consider σ , some k -colouring of G , and some $q \in [k]$ such that $\sigma_v \neq q$. It is easy to see that σ specifies a *connected* subgraph of G which includes v while every vertex in this subgraph is assigned colouring either q or σ_v . The *maximal induced subgraph* of this kind is called “disagreement graph”¹. Figure 1 shows a 3-colouring. The fat lines indicate the disagreement graph specified by using the colour “g”.

It is direct to show that the disagreement graph that is specified by the colouring σ and the colour q is always a connected, bipartite graph whose parts are coloured σ_v and q , respectively.

Definition 3 Assume that σ , a k -colouring of G , and $q \in [k] \setminus \{\sigma_v\}$ define the disagreement graph Q . The k -colouring of G , σ' is called “ q -switching of σ ” if it is derived from σ by switching the colour assignments of the vertices of G that correspond to the two parts of Q .

In Figure 2 we present the “ q -switching” of the colouring in Figure 1. It is direct that for the colouring σ and for some $q \in [k] \setminus \{\sigma_v\}$ there is a unique q -switching of σ . Also, it is straightforward to show that the q -switching of any proper k -colouring of G is a proper colouring, as well ².

Generally the q -switching of a bad colouring is not always a good. However, given some technical conditions which hold with probability $1 - n^{-\Omega(1)}$ over the choices of G , we show the following, *non-trivial*, statement

¹For a more formal definition of “disagreement graph” see in Section 3.1.

² E.g. see proof of Lemma 2




Figure 1: “Disagreement graph”.




Figure 2: “g-switching”.

The distribution of the q -switching of Z , a bad random k -colouring of G , is very close to the distribution of the good random k -colourings of G , when the colour q is chosen uniformly at random from $[k] \setminus \{Z(v)\}$ and k is sufficiently large.

The above fact suggests that we can have the the following *approximation algorithm* for Problem 1 when G is a “typical” instance: Let X be a random colouring G . If X is *good*, then set $X' = X$. If X is a *bad*, then choose at random some $q \in [k] \setminus \{X(v)\}$ and set X' to be equal to the q -switching of X .

Remark. The algorithm in the previous paragraph is exactly the one we refer in the introduction as STEP.

Returning to the approximate random colouring algorithm, we can build upon STEP as follows. First, colour randomly G_0 with some known algorithm. Then, for $i = 0$ to $r - 1$ do the following: If the colouring of G_i is *good*, then consider it as the colouring for G_{i+1} . Otherwise, choose appropriately a random colour q and set as a colouring for the graph G_{i+1} the q -switching of the colouring of G_i .

The above is a concise description of our approximate random colouring algorithm. Clearly it is efficient and accurate only for typical instances of the input graph $G(n, d/n)$, i.e. it has the properties described by Theorem 1 and Theorem 2.

2.1 Some further remarks

To get a better intuition about the algorithm STEP we focus on a case where things go wrong, i.e. consider the following. Let σ be bad colouring of G , i.e. $\sigma_v = \sigma_u$. It is possible that the disagreement graph specified by σ and some colour q to be so large that it contains both v and u . In this case the q -switching of σ is a bad colouring. Clearly in this case STEP fails to generate a good colouring of G . Moreover, it is possible to have good colourings of G that cannot be generated by applying the algorithm STEP to any bad colouring of G . Such colourings constitute *pathological* cases for the algorithm. These pathological cases do not cause big problem as long as they occur rarely, i.e. the fraction of colourings of G that causes such situation is sufficiently small. The occurrences of pathological cases are rare when k is large and v, u are far apart.

3 Problem 1 and α -isomorphism

STEP uses the idea of q -switching so as to achieve a certain kind of mapping between bad and good colourings. *Ideally* this mapping should have the property that, for a bad random colouring of G on the input, the image should be a good random colouring of G . Unfortunately the q -switchings (as implemented by STEP) do not have this property but somehow they *approximate* such mapping. We introduce few notions which capture the essence of these ideas. For the the following definitions in this

section consider a fixed graph G and let Ω be the sets of its proper k -colourings³.

Definition 4 (Isomorphism) We let $\Omega_1, \Omega_2 \subseteq \Omega$. We say that Ω_1 is isomorphic to Ω_2 if and only if there is a bijection $T : \Omega_1 \rightarrow \Omega_2$.

The basic property of isomorphism we need here is contained in the following corollary.

Corollary 1 Assume that we have two isomorphic sets Ω_1 and Ω_2 and let T be a bijection between these two sets. Then, given X_1 , a random member of Ω_1 , the distribution of $T(X_1)$ is the uniform over Ω_2 .

The proof of Corollary 1 appears in Section C.7. The previous definition of isomorphism is standard and generally it expresses a notion of “similarity”. We will need to get a bit further from this, i.e. we introduce a more general notion of “similarity” between sets of colourings which we call α -isomorphism.

Definition 5 (α -isomorphism) We let $\Omega_1, \Omega_2 \subseteq \Omega$ and $\alpha \in [0, 1]$. We say that Ω_1 is α -isomorphic to Ω_2 if there are sets $\Omega'_1 \subseteq \Omega_1$ and $\Omega'_2 \subseteq \Omega_2$ such that

- $|\Omega'_i| \geq (1 - \alpha)|\Omega_i|$, for $i = 1, 2$.
- Ω'_1 and Ω'_2 are isomorphic.

We call (Ω'_1, Ω'_2) as the isomorphic pair of Ω_1 and Ω_2 .

Thus, rather than asking for the whole sets Ω_1 and Ω_2 to be isomorphic, α -isomorphicity requires only sufficiently large subsets from each of Ω_1 and Ω_2 to be isomorphic. The notion of α -function, that follows, is for α -isomorphism the analogous of the bijection for isomorphism.

Definition 6 (α -function) For some $\alpha \in [0, 1]$, let $\Omega_1, \Omega_2 \subseteq \Omega$ be two sets such that Ω_1 is α -isomorphic to Ω_2 with isomorphic pair (Ω'_1, Ω'_2) . Let $h : \Omega'_1 \rightarrow \Omega'_2$ be a bijection. Then the function $H : \Omega_1 \rightarrow [k]^V$ is called α -function if and only if $\forall \sigma \in \Omega'_1$ it holds that $H(\sigma) = h(\sigma)$.

Note that we can be a bit loose on the definition of an α -function when the input σ does not belong to Ω'_1 , i.e. we allow the α -function to take any value in $[k]^V$. Showing that two sets Ω_1 and Ω_2 are α -isomorphic reduces to providing a function which has the properties stated in Definition 6.

Typically, we are given two sets of k -colourings of G , e.g. Ω_1 and Ω_2 , and we will be asked to *devise* an α -function which then suggests that these Ω_1 and Ω_2 are α -isomorphic. The challenge is to devise an α -function H which complies to the following *efficiency rules*: First, given some $\sigma \in \Omega_1$ we want $H(\sigma)$ to have as few different colour assignments from σ as possible, while the vertices with the different colour assignments should be as close to each other as possible. Second, the smaller α and k are, the better. The q -switchings we introduced in the previous section are examples of α -functions between certain sets of colourings of G . The next lemma states the most important property of α -isomorphism and somehow it generalizes Corollary 1.

Lemma 1 Assume that the set Ω_1 is α -isomorphic to Ω_2 , and $H : \Omega_1 \rightarrow [k]^V$ is an α -function. Let z be a random variable distributed uniformly over Ω_1 and let $z' = H(z)$. Denote by ν the uniform distribution over Ω_2 and ν' the distribution of z' . It holds that

$$\|\nu - \nu'\| \leq \alpha.$$

The proof of Lemma 1 appears in the appendix, Section C.1.

³Take k sufficiently large that Ω is non-empty.

3.1 Dealing with Problem 1

In this section we focus on STEP. For clarity reasons we describe the algorithm by assuming that the graph G in Problem 1 is some general fixed graph. A basic part of the presentation involves relating the accuracy of STEP to α -isomorphism between certain sets of k -colourings of G .

Let us introduce some notation. Let Ω denote the set of k -colourings of G and for $c, q \in [k]$ we let $\Omega(c, q) \subseteq \Omega$ denote all the k -colourings of G that assign v and u the colours c, q , respectively. We define formally a disagreement graph as follows:

Definition 7 (Disagreement graph) For $\sigma \in \Omega$ and some $q \in [k] \setminus \{\sigma_v\}$ we let the disagreement graph $Q_{\sigma_v, q} = (V', E')$ be the maximal induced subgraph of G such that

$$V' = \left\{ x \in V \mid \begin{array}{l} \exists \text{ path } w_0, \dots, w_t, \text{ in } G \text{ such that:} \\ w_0 = v, w_t = x, \sigma(w_j) \in \{\sigma_v, q\} \end{array} \right\}.$$

It is important to remember that the disagreement graph is always connected, bipartite and maximal, i.e. for every σ and q , G has no vertex $y \notin Q_{\sigma_v, q}$ which has a neighbour in V' and at the same time $\sigma_y \in \{\sigma_v, q\}$. Furthermore, we define formally the q -switchings as a function $H : \Omega \times [k] \rightarrow [k]^V$, i.e. $H(\sigma, q)$ returns the q -switching of σ .

Function H

Input: $X \in \Omega$ and $q \in [k] \setminus \{X(v)\}$

Set $c = X(v)$.

Set $V_1 = \{w \in Q_{X(v), q} \mid X(w) = c\}$.

Set $V_2 = \{w \in Q_{X(v), q} \mid X(w) = q\}$.

$\forall w \in V_1$ set $X(w) = q$.

$\forall w \in V_2$ set $X(w) = c$.

Output: X

We have reached the point where we have all the definitions we need to describe the algorithm STEP.

STEP

Input: $X \in \Omega$, and k

If X is a *good* colouring of G , then set $Y = X$.

If X is a *bad* colouring of G , then choose q u.a.r. from $[k] \setminus \{X(v)\}$ and set $Y = H(X, q)$.

Output: Y

As far as the accuracy of STEP is concerned we have to show that if X is a bad random k -colouring of G , then $H(X, q)$, as calculated by the algorithm, is distributed sufficiently close to the desired distribution. To this end α -isomorphism comes into use.

For any $c, q \in [k]$ we let $S(c, c) \subseteq \Omega(c, c)$ and $S(q, c) \subseteq \Omega(q, c)$ be defined as follows: The set $S(c, c)$ contains every $\sigma \in \Omega(c, c)$ with the property that the disagreement graph $Q_{\sigma_v, q}$ does not contain the vertex u . Similarly, $S(q, c)$ contains every $\sigma \in \Omega(q, c)$ such that the disagreement graph $Q_{\sigma_v, c}$ does not contain the vertex u . We show that these two sets are isomorphic.

Lemma 2 For any $c, q \in [k]$ with $c \neq q$, it holds that $S(c, c)$ and $S(q, c)$ are isomorphic and the function $H(\cdot, q) : S(c, c) \rightarrow S(q, c)$ is a bijection.

The proof of Lemma 2 appears in the Section C.2. Based on the previous consideration, we provide a general relation between α -isomorphism and the accuracy of STEP. We make the following assumption.

Assumption 1 For some $\alpha \in [0, 1]$ it holds that $\Omega(c, c)$ is α -isomorphic to $\Omega(q, c)$ with α -function $H(\cdot, q)$, for any $c, q \in [k]$ such that $c \neq q$.

Clearly, $(S(c, c), S(q, c))$ is the isomorphic pair of the α -isomorphism between $\Omega(c, c)$ and $\Omega(q, c)$. Assumption 1 imposes an upper bound for the number of *pathological* colourings⁴ of G . It implies that, for any $c, q \in [k]$ all but an α fraction of the colourings in $\Omega(q, c)$ do not have disagreement graph $Q_{q,c}$ which includes both v and u . The same should hold for $\Omega(c, c)$ for disagreement graph $Q_{c,q}$.

Theorem 3 Let ν be the uniform distribution over the good k -colourings of G . Let, also, ν' be the distribution of the output of STEP when the input colouring is distributed uniformly over the k -colourings of G . Under Assumption 1 it holds that

$$||\nu - \nu'|| \leq \alpha,$$

where α is defined in Assumption 1.

The proof of Theorem 3 appears in the appendix, Section C.3. The impact of Assumption 1 to the accuracy of STEP is apparent.

The value of α in Assumption 1 depends on G , k and the function H . The natural way of considering that $\Omega(c, c)$ is α -isomorphic to $\Omega(c, q)$ is mainly as a consequence of the α -function $H(\cdot, q)$. That is, the two sets have this property because we have devised a mapping, the $H(\cdot, q)$, which happens to be an α -function between the two sets. Consequently, someone could device a “better” function, i.e. an α' -function for $\Omega(c, c)$ and $\Omega(c, q)$ such that either $\alpha' < \alpha$, or $\alpha' = \alpha$ but allowing smaller k , or both.

Since the algorithm STEP implements the α -function, the performance of the α -function reflects the performance of the algorithm itself. Clearly, the α -function should be computable in polynomial time.

Lemma 3 For a graph $G = (V, E)$ and some integer k , the time complexity of computing the function $H(\cdot, q)$ is $O(|E|)$.

The proof of Lemma 3 appears in Section C.4.

4 From the algorithm Step to Random Colouring.

Here, we give a general presentation of the approximate random colouring algorithm, which builds upon STEP. We also study properties of the algorithm like time complexity and accuracy. In particular, we study the accuracy of the algorithm under general assumptions about α -isomorphism, as we did in Section 3.1 for STEP. As in the previous cases, the input graph G is considered to be fixed.

First, we extend the notation of the previous section to fit here. For input graph G the algorithm considers the sequence of subgraphs G_0, G_1, \dots, G_r . G_i is derived by deleting from G_{i+1} an edge which we call $\{v_i, u_i\}$. Let Ω_i be the set of k -colourings of G_i . For any $c, q \in [k]$ we let $\Omega_i(c, q)$ be the set of colourings of G_i which assign the colours c and q to the vertices v_i and u_i , respectively.

We proceed by describing the full algorithm in pseudocode. The variable Y_i , below, denotes the k -colouring that the algorithm assigns to the graph G_i .

Random Colouring Algorithm

Input: G, k .
 Compute G_0, G_1, \dots, G_r .
 Compute Y_0 . /* Get a random k -colouring of G_0 .*/
 For $0 \leq i \leq r - 1$ do
 Set Y_{i+1} the output of STEP with input Y_i .

⁴see also discussion in Section 2.1

Output: Y_r .

In the second line the algorithm computes the sequence of subgraphs and in the third it colours randomly G_0 . A detailed description of how can someone construct the sequence of subgraphs and colour randomly G_0 is a graph specific problem. For the case where the input graph is an instance of $G(n, d/n)$, we give a detailed treatment of in the proof of Theorem 1 and Theorem 2⁵. However, using Lemma 3 it is direct to get the following theorem.

Theorem 4 *Under the condition that G_0 can be k -coloured randomly in polynomial time, the random colouring algorithm has polynomial time complexity.*

The next issue we have to investigate is the accuracy of the algorithm. As in Section 3.1 we relate the accuracy of the random k -colouring algorithm with α -isomorphism by using the following assumption.

Assumption 2 *For $i = 0, \dots, r-1$ and some $\alpha \in [0, 1]$ it hold that $\Omega_i(c, c)$ is α -isomorphic to $\Omega_i(q, c)$ and $H(\cdot, q)$ is a α -function, for any $c, q \in [k]$ such that $c \neq q$.*

The α -function H is the same as the one defined in Section 3.1. Let $(S_i(c, c), S_i(q, c))$ be the isomorphic pair of the α -isomorphism between $\Omega_i(c, c)$ and $\Omega_i(q, c)$. The sets $S_i(c, c)$ and $S_i(q, c)$ are defined in the same manner as $S(c, c)$ and $S(q, c)$, in Section 3.1. From Assumption 2, we get the following theorem.

Theorem 5 *Let μ be the uniform distribution over the k -colourings of the input graph G . Let, also, μ' be the distribution of the colourings that is returned by the random colouring algorithm. Under Assumption 2, it holds that*

$$||\mu - \hat{\mu}|| \leq r \cdot \alpha,$$

where r is the maximum index in the sequence G_0, G_1, \dots, G_r .

The proof of Theorem 5 appears in the appendix, Section C.5.

5 Proof sketch for Theorem 1

Due to space limitations, in the remaining pages we give a proof sketch our main result, Theorem 1. That is, we consider the random colouring algorithm with input an instance of $G(n, d/n)$ and we let k be the number of colours. From a technical perspective there are two issues to deal with. The first is how do we construct the sequence of subgraphs. The second is to replace the rather general Assumption 2 about α -isomorphism between colour sets with specific results for the graphs G_0, G_1, \dots, G_r .

When the algorithm constructs the sequence of subgraphs it should take into consideration the previous remark that it is preferable in the graph G_i the vertices v_i and u_i to be at a sufficiently large distance. To see why we need this property we provide the following corollary, which follows directly from previous definitions.

Corollary 2 *Consider some fixed graph G_i and $c, q \in [k]$. The set $\Omega_i(c, c)$ is α -isomorphic to $\Omega_i(q, c)$ with α -function $H(\cdot, q)$ if and only if the following holds: Choose u.a.r. a colouring from $\Omega_i(c, c)$ and let $Q_{c,q}$ be the disagreement graph specified by this colouring and q . It should hold that*

$$a \geq Pr[v_i, u_i \in Q_{c,q} | G_i]. \quad (2)$$

Additionally, the analogous condition should hold for a random colouring of $\Omega_i(q, c)$.

⁵See Section A.

Since we are interested in the minimum possible value for α , we try to minimize the probability term in (2). Clearly, the greater the distance between v_i and u_i the less probable is for $Q_{c,q}$ to include them both and, consequently, the more accurate the random colouring algorithm gets. To this end we use the following lemma to construct the sequence of subgraphs.

Lemma 4 *With probability at least $1 - n^{-2/3}$ we can have the sequence G_0, \dots, G_r satisfying the following two properties.*

1. G_0 consists only of isolated vertices and simple cycles, each of maximum length less than $\frac{\log n}{9 \log d}$.
2. In G_i , the graph distance between v_i and u_i is at least $\frac{\log n}{9 \log d}$.

Additionally it holds that

$$Pr[r \geq (1 + n^{-1/3})dn/2] \leq \exp(-n^{1/4}).$$

In the rest of the analysis of the algorithm we assume that the sequence of subgraphs is such that the distance between v_i and u_i is at least $\gamma \log n$, where $\gamma = (9 \log d)^{-1}$, for $i = 0, \dots, r-1$. Since G_0, \dots, G_r are subgraphs of $G_{n,d/n}$, somehow they are random too. The reader should feel free to assume any arbitrary rule that generates G_i from each instance of $G(n, d/n)$. The only restriction we have is that of the distance between v_i and u_i . Then, we use the following theorem.

Theorem 6 *Take $k \geq (2 + \epsilon)d$, where $\epsilon > 0$ and d is a sufficiently large fixed number. There is β_i such that for any $\alpha \geq \beta_i$ it holds that $\Omega_i(c, c)$ is α -isomorphic to $\Omega_i(c, q)$ and $H(\cdot, q)$ is an α -function while*

$$E[\beta_i] \leq C \cdot n^{-\left(1 + \frac{\epsilon}{45 \log d}\right)},$$

for any $c, q \in [k]$, $C > 0$ is fixed and $i = 0, \dots, r-1$.

The expectation of the quantity β_i is over the graph instances $G(n, d/n)$. Taking $r_0 = (1 + n^{-1/3})dn/2$, Theorem 5 implies that

$$E[||\mu - \hat{\mu}||] \leq E\left[\sum_{i=0}^r \beta_i\right] \leq \sum_{i=0}^{r_0} E[\beta_i | r \leq r_0] + n^2 Pr[r \geq r_0].$$

In the last inequality we use that $\beta_i \in [0, 1]$. It easy to see that $E[\beta_i | r \leq r_0] \leq Pr^{-1}[r \leq r_0] \cdot E[\beta_i]$. Then, Theorem 6 and Lemma 4 suggest that there is fixed $C > 0$ such that $E[||\mu - \hat{\mu}||] \leq C \cdot n^{-\frac{\epsilon}{45 \log d}}$. The theorem follows by applying the Markov inequality.

5.1 Proof sketch for Theorem 6

Consider some fixed instance of G_i and let $c, q \in [k]$ such that $c \neq q$. Choose u.a.r. a colouring from $\Omega_i(c, c)$ and let $Q_{c,q}$ be the disagreement graph that is specified by the colouring we chose and q . Similarly, choose u.a.r. from $\Omega_i(q, c)$ and let $Q_{q,c}$ be the disagreement graph specified by the chosen colouring and c . According to Corollary 2, $\Omega_i(c, c)$ is α -isomorphic to $\Omega(q, c)$ for any $\alpha \geq \beta_i$ such that $\beta_i = Pr[v_i, u_i \in Q_{c,q} | G_i] + Pr[v_i, u_i \in Q_{q,c} | G_i]$. Taking the average over G_i we have

$$E[\beta_i] = Pr[v_i, u_i \in Q_{c,q}] + Pr[v_i, u_i \in Q_{q,c}]. \quad (3)$$

We provide a bound on the probability terms in (3) by using the following proposition.

Proposition 1 *Take $k \geq (2 + \epsilon)d$, for fixed $\epsilon > 0$. Let σ be a k -colouring of G_i that is chosen u.a.r. among $\cup_{c' \in [k]} \Omega_i(c, c')$. For some $q \in [k] \setminus \{c\}$ we let the event $A_i = "v_i \text{ and } u_i \in Q_{\sigma_{v_i}, q}"$. There is a positive constant C such that*

$$Pr[A_i] \leq C \cdot n^{-\left(1 + \frac{\epsilon}{45 \log d}\right)} \quad i = 0, \dots, r.$$

Note that in the above proposition we take a disagreement graph of a colouring chosen u.a.r. among the colourings of an instance of G_i that assign the vertex v_i the colour c (the colouring of u_i is “free”). We show that choosing a u.a.r. a colouring from $\cup_{c' \in [k]} \Omega_i(c, c')$ the probability p for this colouring to be in $\Omega_i(c, c)$ is *constant*, i.e. $p = \Theta(1)$. Then, the law of total probability suggests that $Pr[v_i, u_i \in Q_{c,q}] \leq p \cdot Pr[A_i]$. We work similarly for $Pr[v_i, u_i \in Q_{q,c}]$. The theorem follows.

5.2 Proof sketch for Proposition 1

In the experiment in the statement of Proposition 1, we let $W_i(l)$ denote the number of paths in $Q_{c,q}$ that start at v_i and end at u_i and have length l . By the Markov inequality we get that

$$Pr[A_i] \leq \sum_{l=\gamma \log n}^{\infty} E[W_i(l)], \quad (4)$$

where $\gamma = (9 \log d)^{-1}$. Thus it remains to bound the expectation on the r.h.s.

For a vertex w , we let $deg_i(w)$ be its degree in the graph G_i . Consider the product measure $\mathcal{P}(G_i, k)$ such that each vertex $w \in G_i$ is *disagreeing* with probability $q_w = \frac{1}{k - deg_i(w)}$ and *non-disagreeing* with probability $1 - q_w$. Also, the vertex v_i is disagreeing with probability 1. When $k \leq deg_i(w)$ we set $q_w = 1$. A *path of disagreement* in G_i is any simple path which has all its vertices disagreeing.

Let $\Gamma_i(l)$ denote the number of paths of disagreement between v_i and u_i in G_i , in a configuration chosen according to $\mathcal{P}(G_i, k)$. Through a stochastic order relation we show that for any l it holds

$$E[W_i(l)] \leq E_{\mathcal{P}}[\Gamma_i(l)], \quad (5)$$

where the rightmost expectation is w.r.t. both the measure $\mathcal{P}(G_i, k)$ and G_i . Then taking $k \geq (2 + \epsilon)d$ and sufficiently large d it holds that

$$E_{\mathcal{P}}[\Gamma_i(l)] \leq \Theta(1) \cdot n^{l-1} \left(\frac{d}{n}\right)^l \cdot \left(\frac{1}{(1 + \epsilon/5)d}\right)^l = \Theta(1) \frac{1}{n} (1 + \epsilon/5)^{-l}. \quad (6)$$

The coefficient n^{l-1} comes from the fact that between v_i and u_i there are at most n^{l-1} paths of length l , $(d/n)^l$ is an upper bound for the probability to have a specific path of length l in G_i and the final coefficient is related to the probability for a path of length l to be a “path of disagreement”. The proposition follows by combining (4), (5) and (6).

To get a better picture of why there are not many paths of disagreement when $k \geq (2 + \epsilon)d$ consider q_w the marginal distribution of w in G_i to be disagreeing. For $k \geq (2 + \epsilon)d$ it holds that

$$q_w \leq \frac{1}{(1 + \epsilon/2)d} + Pr[deg_i(w) > (1 + \epsilon/2)d + 1].$$

Clearly, $deg_i(w)$ is dominated by $\mathcal{B}(n, d/n)$. Using Chernoff bounds ⁶ we can show that for any fixed ϵ , the rightmost probability is smaller than $e^{c'd}$, for fixed c' . Then, roughly speaking, we have the following situation: The expected degree of w is at most d . Also, w is disagreeing with probability $q_w < 1/d$, for sufficiently large d . Consequently, for every path of disagreement that enters w the expected number of paths that leave w are $d \cdot q_w < 1$.

⁶Corollary 2.4 in [5].

References

- [1] D. Aldous. *Random walks of finite groups and rapidly mixing Markov chains*. In: Séminaire de Probabilités XVII 1981/82, Springer-Verlag, Berlin. pp. 243-297.
- [2] A. Braunstein, M. Mézard, R. Zecchina: *Survey propagation: an algorithm for satisfiability*. Random Structures and Algorithms 27 (2005) 201226
- [3] M. Dyer, A. Flaxman, A. M. Frieze and E. Vigoda. *Random colouring sparse random graphs with fewer colours than the maximum degree*. Random Struct. and Algorithms 29, pp. 450-465, 2006.
- [4] C. Efthymiou and P. G. Spirakis. *Random sampling of colourings of sparse random graphs with a constant number of colours*. In Theoretical Computer Science 407, pp. 134-154, 2008.
- [5] S. Janson, T. Luczak and A. Ruciński. *Random graphs*. Wiley and Sons, Inc. 2000.
- [6] M. Jerrum and A. Sinclair. *The Markov chain Monte Carlo method: an approach to approximate counting and integration*. In Approximation Algorithms for NP-hard problems (Dorit, Hochbaum ed.) PWS 1996.
- [7] F. Kschischang, B. Frey and H. Loeliger. *Factor Graphs and the Sum-Product Algorithm*. In IEEE Transactions on Information Theory, Vol. 47, No. 2, 2001.
- [8] E. Mossel and A. Sly. *Gibbs Rapidly Samples Colorings of $G_{n,d/n}$* . In journal Probability Theory and Related Fields, Vol. 148, No 1-2, 2010.

Appendix

A Proof of Theorem 1 and Theorem 2

In this section we use results from Section 4 to show Theorem 1 and Theorem 2 when the input of the random colouring algorithm is an instance of $G(n, d/n)$. Essentially there are two issues to deal with, the first is how do we construct the sequence of subgraphs, while the second is to replace the, rather general, Assumptions 2 with more specific results for the colourings of the graphs G_0, G_1, \dots, G_r .

It is easy to construct a sequence of subgraph so as to have G_0 randomly k -coloured in polynomial time (e.g. take it such that G_0 is empty). However, the actual construction of the sequence of subgraphs is a bit more complicated task. It has been remarked very early in this work that in the graph G_i the vertices v_i and u_i are at a sufficiently large distance. To see why we need this property we provide the following corollary, which follows directly from the definitions in the previous sections.

Corollary 3 *Consider some fixed graph G_i and $c, q \in [k]$. The set $\Omega_i(c, c)$ is α -isomorphic to $\Omega_i(q, c)$ with α -function $H(\cdot, q)$ if and only if the following holds: Choose u.a.r. a colouring from $\Omega_i(c, c)$ and let $Q_{c,q}$ be the disagreement graph specified by this colouring and q . It should hold that*

$$a \geq \max_{q \in [k] \setminus \{c\}} \Pr[v_i, u_i \in Q_{c,q} | G_i]. \quad (7)$$

Additionally, the analogous condition should hold for a random colouring of $\Omega_i(q, c)$.

Since we are interested in the minimum possible value for α , we see that the greater the distance between v_i and u_i the less probable is for the disagreement graph to include them both. Thus, the greater the distance between v_i and u_i the more accurate the random colouring algorithm gets. To this end we use the following lemma to construct the sequence of subgraphs of $G_{n,d/n}$.

Lemma 5 *With probability at least $1 - n^{-2/3}$ we can have the sequence G_0, \dots, G_r satisfying the following two properties.*

1. G_0 consists only of isolated vertices and simple cycles, each of maximum length less than $\frac{\log n}{9 \log d}$.
2. In G_i , the graph distance between v_i and u_i is at least $\frac{\log n}{9 \log d}$.

Additionally it holds that

$$\Pr[r \geq (1 + n^{-1/3})dn/2] \leq \exp(-n^{1/4}).$$

The proof of Lemma 5 appears in Section C.6.

In the analysis that follows, we assume that the sequence of subgraphs that is computed by the random colouring algorithm, has the properties stated in Lemma 5. Since G_0, \dots, G_r are subgraphs of $G_{n,d/n}$, somehow they are random too and they depend on d . The reader should feel free to assume any, arbitrary, rule that generates G_i from each instance of $G(n, d/n)$. The only restriction we have is that of the distance between v_i and u_i .

Theorem 7 *Take $k \geq (2 + \epsilon)d$, where $\epsilon > 0$ and d is a sufficiently large fixed number. There is β_i such that for any $\alpha \geq \beta$ it holds that $\Omega_i(c, c)$ is α -isomorphic to $\Omega_i(c, q)$ and $H(\cdot, q)$ is an α -function while*

$$E[\beta_i] \leq \frac{(40 + 8\epsilon)k}{\epsilon} n^{-\left(1 + \frac{\epsilon}{45 \log d}\right)},$$

for any $c, q \in [k]$ and $i = 0, \dots, r$.

Since the graphs G_i are random the corresponding sets Ω_i are random too. The above expectation is taken w.r.t. the random graph G_i , for $i = 0, \dots, r-1$. The proof of Theorem 7 appears in Section B.

Proof of Theorem 1: Using Theorem 5 and Theorem 7 we have that

$$E[||\mu - \hat{\mu}||] \leq E\left[\sum_{i=0}^r \beta_i\right],$$

where the expectation is taken over the instances of the input $G_{n,d/n}$. Noting that $\beta_i \in [0, 1]$, we get

$$E[||\mu - \hat{\mu}||] \leq \sum_{i=0}^{(1+n^{-1/3})dn/2} E[\beta_i | r \leq (1+n^{-1/3})dn/2] + n^2 Pr[r \geq (1+n^{-1/3})dn/2].$$

It is direct that

$$E[\beta_i | r \leq (1+n^{-1/3})dn/2] \leq Pr^{-1}[r \leq (1+n^{-1/3})dn/2] \cdot E[\beta_i] \leq \frac{3(40+8\epsilon)k}{2\epsilon} n^{-(1+\frac{\epsilon}{45\log d})}$$

in the final inequality we used Theorem 7. Combining all the above with Lemma 5, we get that

$$E[||\mu - \hat{\mu}||] \leq C \cdot n^{-\frac{\epsilon}{45\log d}}.$$

for fixed $C > 0$. The theorem follows by applying the Markov inequality. \diamond

Proof of Theorem 2: As we show in the proof of Lemma 5, with probability at least $1 - \exp(-n^{1/4})$, the number of edges of $G_{n,d/n}$ is at most $(1+n^{-1/3})\frac{dn}{2}$. From now on in the proof, assume that we are dealing with a graph with $\Theta(n)$ edges. In this case it is direct that $r = \Theta(n)$, as well.

Since the number of edges is linear, we need $O(n)$ time to find whether some edge belongs to a small cycle, i.e. a cycle of length less than $\frac{\log n}{9\log d}$, or not. This can be done by exploring the structure of the $\frac{\log n}{9\log d}$ -neighbourhood around this edge. Thus, the algorithm requires $O(n^2)$ time to create the sequence of subgraph.

Also, it is clear that we need $O(n)$ time to implement one switching of a colouring. For more details on how this can be done see in the proof of Lemma 3. Since $r = O(n)$, we need $O(n^2)$ time for the all colour switchings in the algorithm.

As far as the random colouring of G_0 is regarded we note the following: Using Dynamic Programming we can compute exactly the number of list colourings of a tree T . In the list colouring problem every vertex $v \in T$ has a set $List(v)$ of valid colours, where $List(v) \subseteq [k]$ and v only receives a colour in $List(v)$. For a tree on l vertices, using dynamic programming we can compute the exact number of list colourings in time lk . For a unicyclic component, i.e. a tree with an extra edge, we can consider all the k^2 colourings of the endpoints of the extra edges and for each of these colourings recurse on the remaining tree. Thus, it is direct to show that we can have a random k -colouring of G_0 in time $O(n)$.

The theorem follows by noting that the construction of the sequence of subgraphs with the desired properties fails with probability at most $n^{-2/3}$. \diamond

B Proof of Theorem 7

So as to prove Theorem 7 we use the following proposition.

Proposition 2 Take $k \geq (2 + \epsilon)d$, where $\epsilon > 0$ and d is a sufficiently large number. Let σ be a k -colouring of G_i that is chosen u.a.r. among $\cup_{c' \in [k]} \Omega_i(c, c')$. For some $q \in [k] \setminus \{c\}$ we let the event $A_i = "v_i \text{ and } u_i \in Q_{\sigma_{v_i, q}}"$. It holds that

$$Pr[A_i] \leq \frac{10 + 2\epsilon}{\epsilon} n^{-\left(1 + \frac{\epsilon}{45 \log d}\right)} \quad i = 0, \dots, r.$$

The reader should remark that since the graph G_i is random, for the probability term $Pr[A_i]$ in the proposition it holds that

$$Pr[A_i] = E[Pr[A_i|G_i]],$$

where the expectation w.r.t. G_i . The proof of Proposition 2 appears in Section B.1

Proof of Theorem 7: Consider, first, a fixed sequence G_i , for $i = 0, 1, \dots, r$. Assume that we choose a k -colouring u.a.r. among $\Omega_i(c, c)$ and let $Q_{c, q}$ be the disagreement graph specified by the chosen k -colouring and q . Let the event $B_i = "v_i, u_i \in Q_{c, q}"$ in the above experiment.

Similarly, assume that we choose u.a.r. a k -colouring from $\Omega_i(q, c)$ and let $Q_{q, c}$ be the disagreement graph specified by the chosen k -colouring and q . Let the event $C_i = "v_i, u_i \in Q_{q, c}"$ in this experiment.

We let $\beta_i = \max\{Pr[B_i|\Omega_i(c, c)], Pr[C_i|\Omega_i(q, c)]\}$. Corollary 3 implies that for any $\alpha \geq \beta_i$ it holds that the set $\Omega_i(c, c)$ is α -isomorphic to $\Omega_i(q, c)$ with α -function $H(\cdot, q)$. Also, it is straightforward that

$$E[\beta_i] \leq Pr[B_i] + Pr[C_i].$$

The above expectation is taken w.r.t. to the instances G_i .

Assume that we choose u.a.r. a member of a fixed instance of $\cup_{c' \in [k]} \Omega_i(c, c')$ and we denote with E_i the event that the chosen colouring belongs to $\Omega_i(c, c)$. Also let

$$p = \sum_G Pr[E_i|G] \cdot \mathcal{D}_i[G],$$

where, for a fixed graph G , $Pr[E_i|G]$ is equal to the probability to have the event E_i when the sets of k -colourings are specified by the graph G . $\mathcal{D}_i(G)$ is equal to the probability that an instance of G_i is the graph G . Applying the law of total probability we get that

$$\begin{aligned} Pr[A_i] &= Pr[A_i|E_i]Pr[E_i] + Pr[A_i|E_i^c]Pr[E_i^c] \\ &\geq Pr[A_i|E_i]Pr[E_i] = Pr[B_i] \cdot p. \end{aligned}$$

Thus, it holds that

$$Pr[B_i] \leq p^{-1}Pr[A_i].$$

Since we have the value of $Pr[A_i]$ from Proposition 2, we only need to compute a lower bound for the probability p . For a fixed graph G , let μ_G denote the Gibbs distribution of the k -colourings of G . Also let μ_i be defined as follows:

$$\mu_i(\sigma) = \sum_G \mu_G(\sigma) \mathcal{D}_i(G) \quad \forall \sigma \in [k]^V.$$

We use the following claim to compute bounds for p .

Claim 1 Taking $k \geq (2 + \epsilon)d$, where $\epsilon > 0$ is fixed and d is a sufficiently large number, it holds that

$$\max_{\sigma \in \Omega_i} \|\mu_i(\cdot|\sigma_v) - \mu_i(\cdot)\|_u \leq n^{-1} \quad i = 0, \dots, r.$$

It is easy to show that under μ_i the marginal distribution of the colour assignment of the vertex u is the uniform over the set $[k]$. The above claim suggests that for $k \geq (2 + \epsilon)d$ it holds that

$$\left| p - \frac{1}{k} \right| \leq n^{-1}.$$

Thus we get that

$$Pr[B_i] \leq 2kPr[A_i] \leq \frac{(20 + 4\epsilon)k}{\epsilon} n^{-\left(1 + \frac{\epsilon}{45 \log d}\right)}.$$

Using the same arguments we, also, get that

$$Pr[C_i] \leq \frac{(20 + 4\epsilon)k}{\epsilon} n^{-\left(1 + \frac{\epsilon}{45 \log d}\right)}.$$

The theorem follows. \diamond

Proof of Claim 1: First assume that we have a fixed G_i , i.e. the set of colourings is fixed. Let X_i be distributed uniformly over $\cup_{c' \in [k]} \Omega_i(c, c')$ and let Z_i be distributed uniformly over Ω_i . We couple these two variables. The coupling is done as follows. Choose u.a.r. a colour from $[k]$, and set $Z_i(v_i)$ equal to this colour, e.g. let $Z_i(v_i) = q$. We have two cases.

If $q = c$, then we can have an identical coupling between Z_i and X_i . Otherwise, i.e. if $Z_i(v) \neq X_i(v_i)$, we can set $Z_i = H(X_i, q)$.

Claim 2 *In the later case, i.e. when $Z_i = H(X_i, q)$, Z_i is distributed uniformly over the colouring of G_i that assign the vertex v_i the colour q .*

The proof of the claim appear after the end of this proof.

Thus, in the case where $Z_i(v) \neq X_i(v)$ and we set $Z_i = H(X_i, q)$ it is direct to see that $Z_i(u_i) \neq X_i(u_i)$ if and only if the event A_i (as defined in the statement of Proposition 2) holds. Thus we get that

$$Pr[X_i(u_i) \neq Z_i(u_i) | G_i] \leq Pr[A_i | G_i]$$

From the above relation and Proposition 2 we get that

$$Pr[X_i(u_i) \neq Z_i(u_i)] \leq Pr[A_i] \leq n^{-1},$$

The claim follows by using the Coupling Lemma. \diamond

Proof of Claim 2: We remind the reader that X_i is distributed uniformly at random among the k -colourings of G_i that assign the vertex v_i the colour c . It suffice to show that the sets $\Omega_c = \cup_{c' \in [k]} \Omega_i(c, c')$ and $\Omega_q = \cup_{c' \in [k]} \Omega_i(q, c')$ are isomorphic with bijection $H(\cdot, q) : \Omega_c \rightarrow \Omega_q$. The arguments we need to show this are the same as those we use in the proof of Lemma 2.

I.e. first we need to show that for any $\sigma \in \Omega_c$ it holds that $H(\sigma, q)$ is a proper colouring of G_i . Clearly this holds (see the first two paragraphs of the proof of Lemma 2 in section C.2.) Second we need to show that the mapping $H(\cdot, q) : \Omega_c \rightarrow \Omega_q$ is *surjective*, i.e. for any $\sigma \in \Omega_q$ there is a $\sigma' \in \Omega_c$ such that $\sigma = H(\sigma', q)$. It is direct to see that such σ' exists, moreover, it holds that $\sigma' = H(\sigma, c)$. Finally, we need to show that $H(\cdot, q)$ is *one-to-one*, i.e. there are no two $\sigma_1, \sigma_2 \in \Omega_c$ such that $H(\sigma_1, q) = H(\sigma_2, q)$. Using arguments similar to for the *surjective* case it is direct to see that there cannot be such a pair of colourings. The claim follows. \diamond

B.1 Proof of Proposition 2

Consider the probability distribution $\mathcal{L}(G_i, k)$ (or $\mathcal{L}_{G_i, k}$) induced by the following experiment. We have a graph G_i and we choose a k -colouring σ u.a.r. from $\cup_{c' \in [k]} \Omega_i(c, c')$, where $c \in [k]$. Choose u.a.r. a colour from $[k] \setminus \{c\}$, let q be that colour. Create the graph of disagreement $Q_{c, q}$. If $w \in Q_{c, q}$, then w is “disagreeing” otherwise it is “non-disagreeing”. By definition v_i is always in the disagreement graph.

For a vertex w , we denote with $\deg_i(w)$ its degree in the graph G_i . Consider, also, the product measure $\mathcal{P}(G_i, k)$ such that each vertex $w \in G_i$ is *disagreeing* with probability $q_w = \frac{1}{k - \deg_i(w)}$ and *non-disagreeing* with probability $1 - q_w$. Also, the vertex v_i is disagreeing with probability 1. When $k \leq \deg_i(w)$ we set $q_w = 1$.

A *path of disagreement* in G_i is any simple path which has all its vertices disagreeing. The measure $\mathcal{P}(G_i, k)$ will turn out to be very useful because it dominates $\mathcal{L}(G_i, k)$ in the following sense.

Lemma 6 *Let $M = x_1, x_2, \dots, x_l$ be a path in G_i such that $v_i = x_1$. Let the event $E = "M \text{ is a path of disagreement}"$. It holds that*

$$\mathcal{L}_{G_i, k}[E] \leq \mathcal{P}_{G_i, k}[E].$$

Proof: Let the event $E_i = "x_i \text{ is disagreeing}"$, for $i \leq l$, obviously $E = \bigcap_{j=1}^l E_j$. It is direct that

$$\mathcal{L}_{G_i, k}[E] = \mathcal{L}[E_1] \prod_{j=2}^l \mathcal{L}_{G_i, k}[E_j | \cap_{s=1}^{j-1} E_s].$$

The path of disagreement is specified by a random colouring from $\cup_{c' \in [k]} \Omega_i(c, c')$, call this random colouring X . Let also N_j be the vertices which are adjacent to the vertex x_j . W.l.o.g. assume that $k > \deg_i(x_j)$, for $j \leq l$. Clearly it holds that

$$\mathcal{L}_{G_i, k}[E_j | \cap_{s=1}^{j-1} E_s] \leq \max_{\sigma \in \cup_{c' \in [k]} \Omega_i(c, c')} \mathcal{L}_{G_i, k}[E_j | X(N_j) = \sigma_{N_j}] \leq \frac{1}{k - \deg_i(x_j)}.$$

Thus

$$\mathcal{L}_{G_i, k}[E] \leq \prod_{j=1}^l \frac{1}{k - \deg_i(x_j)} \leq \mathcal{P}_{G_i, k}[E].$$

The lemma follows. ◇

The following corollary is straightforward.

Corollary 4 *Let $M = x_1, x_2, \dots, x_l$ be a path in G_i . Let the event $E = "M \text{ is a path of disagreement}"$. It holds that*

$$\mathcal{P}_{G_i, k}[E] \leq \mathcal{P}_{G_{i+1}, k}[E].$$

We, also, need the following lemma for the proof of Proposition 2.

Lemma 7 *Consider the product measure $\mathcal{P}(G_{n, d/n}, k)$, for $k \geq (2 + \epsilon)d$, for fixed $\epsilon > 0$. Let π be a permutation of $l + 1$ vertices of $G_{n, d/n}$, for $0 \leq l \leq \Theta(\log^2 n)$. There exists $d_0(\epsilon)$, such that for $d > d_0(\epsilon)$ it holds that*

$$\mathcal{P}_{G_{n, d/n}, k}[\pi \text{ is a path of disagreement}] \leq \left(\frac{d}{n}\right)^l \cdot \left(\left(\frac{1}{(1 + \epsilon/4)d} + 3n^{-0.95} \right)^l + 2n^{-\log^4 n} \right).$$

Proof: Call π the path that corresponds to the permutation π , e.g. $\pi = (x_1, \dots, x_{l+1})$. Let Γ be an indicator variable such that $\Gamma = 1$ if π is a path of disagreement and $\Gamma = 0$, otherwise. Let, also, I_π be the event that there exists the path (x_1, \dots, x_{l+1}) in $G_{n,d/n}$. It holds that

$$E_{\mathcal{P}}[\Gamma] = \left(\frac{d}{n}\right)^l \cdot E_{\mathcal{P}}[\Gamma|I_\pi].$$

Let Q_π denote the event that the vertices in π have degree less than $\log^6 n$. Using Chernoff bounds it is easy to show that $Pr[Q_\pi|I_\pi] \geq 1 - n^{-\log^4(n)}$. Also, it holds that

$$\begin{aligned} E_{\mathcal{P}}[\Gamma|I_\pi] &= E_{\mathcal{P}}[\Gamma|I_\pi, Q_\pi]Pr[Q_\pi|I_\pi] + E_{\mathcal{P}}[\Gamma|I_\pi, \bar{Q}_\pi]Pr[\bar{Q}_\pi|I_\pi] \\ &\leq E_{\mathcal{P}}[\Gamma|I_\pi, Q_\pi] + n^{-\log^4(n)}. \end{aligned}$$

It suffice to show that for $0 \leq l \leq \Theta(\log^2 n)$ and sufficiently large n it holds that

$$E_{\mathcal{P}}[\Gamma|I_\pi, Q_\pi] \leq \left(\frac{1}{(1 + \epsilon/4)d} + 3n^{-0.95}\right)^l. \quad (8)$$

We show (8) by induction on l . Clearly for $l = 0$ the inequality in (8) is true. Assuming that (8) holds for $l = l_0$, we will show that it holds for $l = l_0 + 1$, as well.

For a vertex w , we let $D(w)$ denote the event that this vertex is disagreeing. Given that all vertices in $\{x_1, \dots, x_{l_0}\}$ are disagreeing we let $deg_{out}(x_i)$ be the number of vertices in $V \setminus \{x_1, \dots, x_{l_0}\}$ that are adjacent to x_i , for $1 \leq i \leq l_0$. If $deg_{out}(x_i) = t$, then all the possible subsets of $V \setminus \{x_1, \dots, x_{l_0}\}$ with cardinality t are equiprobably adjacent to x_i . This implies that

$$Pr[x_{l_0+1} \text{ is adjacent to } x_i] = \frac{E[deg_{out}(x_i)]}{n - l_0} \quad \text{for } 0 \leq i \leq l_0 - 1.$$

Let $deg_{in}(x_{l_0+1})$ be the number of neighbours of x_{l_0+1} in $\{x_1, \dots, x_{l_0-1}\}$. By the linearity of expectation we have

$$E[deg_{in}(x_{l_0+1})|I_\pi, Q_\pi] \leq \frac{l_0}{n - l_0} E[deg_{out}(x_i)|I_\pi, Q_\pi] \leq n^{-0.95}. \quad (9)$$

We make the simplifying assumption that if the vertex x_{l_0+1} is adjacent to any vertex in $\{x_1, \dots, x_{l_0-1}\}$, then it is disagreeing, regardless of the number of adjacent vertices outside the path. By (9) and the Markov inequality, we get that

$$Pr[deg_{in}(x_{l_0+1}) > 0|I_\pi, Q_\pi] \leq E[deg_{in}(x_{l_0+1})|I_\pi, Q_\pi] \leq n^{-0.95}.$$

We denote with E the event that “ (x_1, \dots, x_{l_0}) is a path of disagreement, $deg_{in}(x_{l_0+1}) = 0$, the edge $\{x_{l_0}, x_{l_0+1}\}$ appears in $G_{n,d/n}$ and the event Q_π holds”. It is easy to show that $Pr[E] \geq 1 - 2n^{-0.95}$. It holds that

$$\begin{aligned} Pr[D(x_{l_0+1})|E] &\leq \sum_{j=0}^n Pr[D(x_{l_0+1})|E, deg_{out}(x_{l_0+1}) = j] Pr[deg_{out}(x_{l_0+1}) = j|E] \\ &\leq (1 + 3n^{-0.95}) \sum_{j=0}^{k-1} \frac{1}{k-j} \binom{n}{j} (d/n)^j (1 - d/n)^{n-j} \\ &\quad + (1 + 3n^{-0.95}) \sum_{j=k}^n \binom{n}{j} (d/n)^j (1 - d/n)^{n-j} \\ &\leq q(k, d) + 3n^{-0.95} \end{aligned}$$

where

$$q(k, d) = \sum_{j=0}^{k-1} \frac{1}{k-j} \binom{n}{j} (d/n)^j (1-d/n)^{n-j} + \sum_{j=k}^n \binom{n}{j} (d/n)^j (1-d/n)^{n-j}.$$

The following inequalities are straightforward.

$$\begin{aligned} q(k, d) &\leq \sum_{j=0}^{k/2} \frac{1}{k-j} \binom{n}{j} (d/n)^j (1-d/n)^{n-j} + \sum_{j=k/2+1}^n \binom{n}{j} (d/n)^j (1-d/n)^{n-j} \\ &\leq \frac{2}{k} + \Pr[B(n, d/n) \geq (1 + \epsilon/2)d + 1]. \end{aligned}$$

Using Chernoff bounds, i.e. Corollary 2.4 from [5] we get that

$$\Pr[B(n, d/n) \geq (1 + \epsilon/2)d + 1] \leq \exp(-c'd)$$

where $c' = \log \epsilon - 1 + \frac{1}{1+\epsilon} > 0$. It is clear that taking $k \geq (2 + \epsilon)d$ for fixed $\epsilon > 0$, there is sufficiently large $d_0(\epsilon)$ such that for $d > d_0(\epsilon)$ it holds that

$$q(k, d) \leq \frac{2 + 2/d}{k} \leq \frac{1}{(1 + \epsilon/4)d}.$$

The lemma follows. \diamond

Proof of Proposition 2: Let the event $B = “v_i \text{ and } u_i \text{ are connected through a path of disagreement of length at most } \log^2 n”$. Also, let the event $C = “v_i \text{ and } u_i \text{ are connected through a path of length greater than } \log^2 n”$. Clearly it holds that

$$\mathcal{L}_{G_i, k}[A] \leq \mathcal{L}_{G_i, k}[B] + \mathcal{L}_{G_i, k}[C],$$

where $\mathcal{L}_{G_i, k}$ is the probability distribution we defined at the begining of this section. When there is no danger of confusion we drop the subscript G_i, k . The proposition will follow by calculating the probabilities $\mathcal{L}[B]$ and $\mathcal{L}[C]$.

Consider an enumeration of all the permutations of l vertices in G_i with first the vertex v_i and last the vertex u_i . Let $\pi_0(l), \pi_1(l), \dots$ be the permutations in the order they appear in the enumeration. Let $\Gamma_j(l)$ be the random variable such that

$$\Gamma_j(l) = \begin{cases} 1 & \text{the path that corresponds to } \pi_i(l) \text{ is a path of disagreement} \\ 0 & \text{otherwise.} \end{cases}$$

Let, also, $\Gamma(l) = \sum_j \Gamma_j(l)$. It is easy to see that the number of sumads in the previous sum are at most n^{l-1} . Towards computing $\mathcal{L}(C)$, we need to calculate the following expectation

$$E_{\mathcal{L}} \left[\sum_{l=l_0}^{\log^2 n} \Gamma(l) \right],$$

where $l_0 = \frac{\log n}{9 \log d}$. However, we have to take into consideration that we have conditioned that v_i and u_i are at distance at least $\frac{\log n}{9 \log d}$. To this end, it is direct to show that if Z the number of paths of length at most $\frac{\log n}{9 \log d} - 1$ between two vertices of G_i , then

$$E[Z] \leq \sum_{l \leq \frac{\log n}{9 \log d} - 1} n^{l-1} \left(\frac{d}{n} \right)^l \leq n^{-9/10}.$$

Thus, letting \hat{p} be the probability of the event that two vertices are at distance is at least $\frac{\log n}{9 \log d}$ the Markov inequality suggests that $\hat{p} \geq 1 - n^{-9/10}$. Using Lemma 6, Lemma 7 and Corollary 4 we get that

$$\begin{aligned} E_{\mathcal{L}} \left[\sum_{l=l_0}^{\log^2 n} \Gamma(l) \right] &\leq \hat{p}^{-1} \sum_{l=l_0}^{\log^2 n} n^{l-1} \left(\frac{d}{n} \right)^l \left(\left(\frac{1}{(1+\epsilon/4)d} + 3n^{-0.95} \right)^l + 2n^{-\log^4 n} \right) \\ &\leq \frac{1}{n\hat{p}} \sum_{l=l_0}^{\log^2 n} \left(((1+\epsilon/4)^{-1} + 3dn^{-0.95})^l + 2d^l n^{-\log^4 n} \right). \end{aligned}$$

Note that $d^l n^{-\log^4 n} = O(n^{-\log^4 n})$, for $l = O(\log^2 n)$. Thus, for sufficiently large n and d we get that

$$\begin{aligned} E_{\mathcal{L}} \left[\sum_{l=l_0}^{\log^2 n} \Gamma(l) \right] &\leq \sum_{l=l_0}^{\log^2 n} \frac{3}{2n} \left(1 + \frac{\epsilon}{5} \right)^{-l} \\ &\leq \frac{3}{2n} \left(1 + \frac{\epsilon}{5} \right)^{-l_0} \frac{1}{1 - (1 + \epsilon/5)^{-1}} \leq \frac{15 + 3\epsilon}{2\epsilon} n^{-\left(1 + \frac{\epsilon}{45 \log d}\right)}. \end{aligned}$$

Using the Markov inequality we get that

$$\mathcal{L}[B] \leq E_{\mathcal{L}} \left[\sum_{l \geq l_0}^{\log^2 n} \Gamma(l) \right] \leq \frac{15 + 3\epsilon}{2\epsilon} n^{-\left(1 + \frac{t}{45 \log d}\right)}. \quad (10)$$

Let $P(l)$ be the number of paths of disagreement between v_i and *any* vertex of G_i , that have length l . It is direct that

$$\mathcal{L}[C] \leq \Pr [P(\log^2 n) > 0].$$

The above inequality follows by noting that so as to have a path of disagreement connecting v_i and u_i which has length at least l , we should have some path of disagreement of length l leaving v_i . Using Markov's inequality we get that

$$\begin{aligned} \Pr [P(\log^2 n) > 0] &\leq E_{\mathcal{L}} [P(\log^2 n)] \\ &\leq \hat{p}^{-1} n^{\log^2 n - 1} \left(\frac{d}{n} \right)^{\log^2 n} \left(\left(\frac{1}{(1+\epsilon/4)d} + 3n^{-0.95} \right)^{\log^2 n} + 2n^{-\log^4 n} \right) \\ &\leq \frac{1}{\hat{p}n} (1 + \epsilon/5)^{-\log^2 n} = \Theta \left(n^{-\frac{\epsilon}{10} \log n} \right). \end{aligned}$$

The proposition follows. \diamond

C Proofs

C.1 Lemma 1

Proof: Let x be a r.v. distributed as in ν . The proof of this lemma is going to be made by coupling x and z' . In particular, we show that there is a coupling of x and z' such that

$$\Pr[x \neq z'] \leq \alpha.$$

Then the lemma will follow by using Coupling Lemma [1]. Let (Ω'_1, Ω'_2) be the isomorphic pair of the α -isomorphism between Ω_1 and Ω_2 . Observe that $|\Omega'_i| \geq (1 - \alpha)|\Omega_i|$, for $i = 1, 2$. Also, it holds that

$$Pr[z' = \sigma | z \in \Omega'_1] = \frac{1}{|\Omega'_2|} \quad \forall \sigma \in \Omega'_2.$$

Note that when we restrict the input of the α -function H only to members of Ω'_1 , then H is by definition a bijection between the sets Ω'_1 and Ω'_2 . The above equality then follows by using Corollary 1 and noting that conditional on the fact that $z \in \Omega'_1$, z is distributed uniformly over Ω'_1 . Also it is easy to get that

$$Pr[x = \sigma | x \in \Omega'_2] = \frac{1}{|\Omega'_2|} \quad \forall \sigma \in \Omega'_2.$$

Let

$$p = \min\{Pr[x \in \Omega'_2], Pr[z \in \Omega'_1]\} \geq 1 - \alpha. \quad (11)$$

The above inequality follows from the assumption that Ω_1 and Ω_2 are α -isomorphic.

It is clear that we can have a coupling between x , z and z' such that the event $E = "z \in \Omega'_1 \text{ and } x \in \Omega'_2"$ holds with probability p (see (11)). In this coupling, if the event E holds, then x and y are distributed uniformly over Ω'_2 and Ω'_1 , respectively. This means that we can make an extra arrangement such that when E holds to have $x = H(z)$, as well. Since $z' = H(z)$, it is direct that when the event E holds we, also, have that $x = z'$. We conclude that in the above coupling it holds that $Pr[x = z'] \geq Pr[E]$. Thus,

$$Pr[x \neq z'] \leq 1 - Pr[E] = 1 - p = \alpha.$$

The lemma follows. \diamond

C.2 Lemma 2

Proof: First we are going to show that for any $\sigma \in S(c, c)$, it holds that $H(\sigma, q)$ is a proper colouring of G . Assume the contrary, i.e. that there is $\sigma \in S(c, c)$ such that $H(\sigma, q)$ is a non-proper colouring, i.e. there is a monochromatic edge e . Let $Q_{\sigma_v, q}$ be the disagreement graph specified by σ and q . It is direct that the monochromatic edge is either incident to two vertices in $Q_{\sigma_v, q}$ or to some vertex in $Q_{\sigma_v, q}$ and some vertex outside the disagreement graph.

It is direct that $H(\sigma, q)$ does not cause any monochromatic edge between two vertices in $Q_{\sigma_v, q}$. To see this, note that the disagreement graph is bipartite and σ specifies exactly one colour for each part of the graph, while $H(\sigma, q)$ switches the colours of the two parts. On the other hand, $H(\sigma, q)$ cannot cause any monochromatic edge between a vertex in $Q_{\sigma_v, q}$ and some vertex outside the disagreement graph. This follows by the fact that the disagreement graph is maximal. Thus, there is no edge w outside $Q_{\sigma_v, q}$ such that $\sigma_w \in \{q, \sigma_v\}$ while at the same time w is adjacent to some vertex in $Q_{\sigma_v, q}$.

Also, it is direct to show that for any $\sigma \in S(c, c)$, it holds that $H(\sigma, q) \in S(q, c)$. This follows by the definition of the sets $S(c, c)$ and $S(c, q)$. It remains to show that $H(\cdot, q) : S(c, c) \rightarrow S(q, c)$ is a bijection.

We show that $H(\cdot, q)$ has range the set $S(q, c)$, i.e. it is *surjective* map, ie. for any $\sigma \in S(q, c)$ there is $\sigma' \in S(c, c)$ such that $\sigma = H(\sigma', q)$. It is direct to see that such σ' exists, moreover, it holds $\sigma' = H(\sigma, c)$.

Finally, we need to show that $H(\cdot, q)$ is *one-to-one*, i.e. there are no two $\sigma_1, \sigma_2 \in S(c, c)$ such that $H(\sigma_1, q) = H(\sigma_2, q)$. Using arguments similar to those in the previous paragraph it is direct to see that there cannot be such a pair of colourings.

Thus, since $H(\cdot, q) : S(c, c) \rightarrow S(q, c)$ is surjective and one-to-one it is a bijection. The lemma follows. \diamond

C.3 Theorem 3

Proof: Let X be the input of STEP, i.e. a random k -colouring of G . Let Y be equal to the colouring that is returned by the algorithm. Also, let Z be a random variable distributed as in ν . The proof of the theorem is going to be made by coupling Z and Y and by showing that in this coupling it holds that

$$Pr[Z \neq Y] \leq \alpha.$$

The reader should observe that for any $q, c \in [k]$ such that $c \neq q$, it holds that

$$Pr[Z(v) = q | Z(u) = c] = Pr[X(v) = q | X(u) = c, X(v) \neq c] = \frac{1}{k-1} \quad (12)$$

and

$$Pr[X(v) = X(u) = c | X \text{ is bad}] = \frac{1}{k}, \quad (13)$$

due to symmetry. Also, it is direct to show that

$$Pr[Y(v) = q | X(u) = c] = \frac{1}{k-1} \quad (14)$$

for every $q \in [k] \setminus \{c\}$. Now we are going to construct the coupling. We need to involve the variable X , the input of STEP, in this coupling. First, set $Z(u) = X(u)$ and then set $Z(v) = Y(v)$. Using the above observations it is straightforward to show that $Z(u)$ and $Z(v)$ are set, respectively, according to the appropriate distribution (due to (12) and (14)).

We reveal the values of $X(v)$, $X(u)$ and $Y(v)$. By the above coupling we also have the values of $Z(v)$ and $Z(u)$. We consider two cases, depending on whether X is a good or a bad colouring.

If $X(v) \neq X(u)$, i.e. X is good, then we have $X = Y$ and we can set directly $X = Z$. Thus, for the coupling it holds

$$Pr[Y \neq Z | X \text{ is good}] = 0.$$

If $X(u) = X(v)$, then w.l.o.g. we can assume $X(u) = X(v) = c$, for some $c \in [k]$. In this case, we choose whether $X \in S(c, c)$ or not. For this choice, the Assumption 1 suggests that

$$Pr[X \in S(c, c) | X(u) = X(v) = c] \geq 1 - \alpha.$$

Similarly for Z , assume that $Z(u) = c$ and $Z(v) = q$, with $c \neq q$. Again Assumption 1 suggests that

$$Pr[Z \in S(q, c) | X(u) = X(v) = c] \geq 1 - \alpha.$$

Let the event $E = "X \in S(c, c) \text{ and } Z \in S(q, c)"$. Having set $X(v)$, $X(u)$, $Z(v)$, $Z(u)$, $Y(v)$, the two previous inequalities suggest that we can couple X and Z such that the probability of the event E to occur is at least $1 - \alpha$.

Claim 3 *Conditional on the event E , Y is distributed uniformly over $S(q, c)$.*

Conditional on the event E , it is easy to observe that Z is, also, distributed uniformly over $S(q, c)$. This observation and Claim 3 suggest that

$$Pr[Z \neq Y | E] = 0.$$

Gathering all the above together and applying the law of total probability we get the following for the coupling:

$$Pr[Z \neq Y] \leq Pr[Z \neq Y | X \text{ is good}] + Pr[Z \neq Y | E] + Pr[\bar{E}] \leq \alpha.$$

The theorem follows. \diamond

Proof of Claim 3: Conditional on the event E , the random variable X is distributed uniformly over $S(c, c)$. Note that $S(c, c)$ and $S(q, c)$ are isomorphic, due to Lemma 2. The same lemma suggests that we can have a bijection between the two isomorphic sets by taking $H(\cdot, q)$ and restricting its input only to colourings in $S(c, c)$. Thus, since X is distributed uniformly over $S(c, c)$, $H(X, q) = Y$ is distributed uniformly over $S(q, c)$, by Corollary 1. the claim follows. \diamond

C.4 Lemma 3

Proof: Note that the time complexity of computing the value of $H(\sigma, q)$ is dominated by the time we need to reveal the disagreement graph $Q_{\sigma_v, q}$, for some $q \in [k]$. We show that we need $O(|E|)$ steps to reveal the disagreement graph $Q_{\sigma_v, q}$.

We can reveal the graph $Q_{\sigma_v, q}$ in steps $j = 0, \dots, |E|$, where E is the set of edges of G . At step 0 the disagreement graph $Q_{\sigma_v, q}(0)$ contains only the vertex v . Given the graph $Q_{\sigma_v, q}(j)$ we construct $Q_{\sigma_v, q}(j+1)$ as follows: Pick some edge which is incident to a vertex in $Q_{\sigma_v, q}(j)$. If the other end of this edge is incident to a vertex outside $Q_{\sigma_v, q}(j)$ that is coloured either σ_v or q then we get $Q_{\sigma_v, q}(j+1)$ by inserting this edge and the vertex into $Q_{\sigma_v, q}(j)$. Otherwise $Q_{\sigma_v, q}(j+1)$ is the same as $Q_{\sigma_v, q}(j)$. We never pick the same edge twice.

It is direct to show that in the above procedure it holds that $Q_{\sigma_v, q} = Q_{\sigma_v, q}(|E|)$. Thus. the time complexity of a q -switching of a given colouring of G is $O(|E|)$. \diamond

C.5 Theorem 5

Proof: Let X_i be a random variable which is distributed uniformly over Ω_i , $i = 0, \dots, r$. It suffices to provide a coupling of X_r and Y_r , such that

$$Pr[X_r \neq Y_r] \leq r \cdot \alpha.$$

Working as in the proof of Theorem 5 we get the following: There is a coupling of X_i, X_{i+1} such that for the event $E_i = "X_i \text{ is good or there are } c, q \in [k] \text{ such that } X_i \in S_i(c, c) \text{ and } X_{i+1} \in S(q, c)"$ it holds that

$$Pr[E_i] \geq 1 - \alpha.$$

Now consider the random variables $Z = (X_0, X_1, \dots, X_{r-1})$ and $Z' = (X_1, X_2, \dots, X_r)$ and $W = (Y_1, \dots, Y_r)$. Consider, also, the event $E = \bigcap_{i=0}^{r-1} E_i$, where E_i is the event defined above. All the above discussion suggests two facts: First, there is a coupling between Z, Z' and W such that

$$Pr[E] \geq 1 - r\alpha.$$

Second, if in this coupling the event E occurs we can have $Z' = W$, i.e. $X_i = Y_i$, for $i = 1, \dots, r$. To see thus, consider the following: If the event E_i occurs we can have either $X_i = X_{i+1}$ or $X_{i+1} = H(X_i, q)$ for appropriate q . When E occurs, we have this property for all $i = 0, \dots, r-1$. A direct inductive argument implies $Z' = W$. The theorem follows by noting that

$$Pr[Y_r \neq X_r] \leq 1 - Pr[E].$$

\diamond

C.6 Lemma 5

Proof: For (1) it suffice to show that with probability at least $1 - n^{-2/3}$ all the cycles of length less than $\frac{\log n}{9 \log d}$ in $G_{n,d/n}$ do not share edges with each other. Let $\gamma = (9 \log(d))^{-1}$. Assume the opposite, there are at least two cycles, each of length at least $\gamma \log n$ that intersect with each other. Then, there must exist a subgraph of $G_{n,d/n}$ that contains at most $2\gamma \log n$ vertices while the number of edges exceeds by 1, or more, the number of vertices.

Let D be the event that in $G_{n,d/n}$ there exists a set of r vertices which have $r + 1$ edges between them. For $r \leq 2\gamma \log n$ we have the following:

$$\begin{aligned} Pr[D] &\leq \sum_{r=1}^{\gamma \log n} \binom{n}{r} \binom{\binom{r}{2}}{r+1} (d/n)^{r+1} (1 - d/n)^{\binom{r}{2} - (r+1)} \\ &\leq \sum_{r=1}^{\gamma \log n} \left(\frac{ne}{r}\right)^r \left(\frac{r^2 e}{2(r+1)}\right)^{r+1} (d/n)^{r+1} \leq \frac{e \cdot d}{2n} \sum_{r=1}^{\gamma \log n} \left(\frac{e^2 d}{2}\right)^r \\ &\leq \frac{C}{n} \left(\frac{e^2 d}{2}\right)^{2\gamma \log n}. \end{aligned}$$

Having $2\gamma \cdot \log(e^2 d/2) < 1$, the quantity in the r.h.s. of the last inequality is $o(1)$, in particular it is of order $\Theta(n^{\gamma \log(e^2 d/2) - 1})$. Thus, for $\gamma = (9 \log d)^{-1}$ there is no connected component that contains two cycles with probability at least $1 - 2n^{-2/3}$.

If we include in G_0 all the edges that belong to small cycles, i.e. of length less than $\frac{\log n}{9 \log d}$ then it is straightforward that (2) holds.

For (3), we let $E(G_{n,d/n})$ be the number of edges in $G_{n,d/n}$. Using standard probabilistic tools, i.e. Chernoff bounds, it is direct to get that

$$Pr \left[E(G_{n,d/n}) \geq (1 + n^{-1/3}) \frac{dn}{2} \right] \leq \exp \left(-n^{1/4} \right).$$

It is direct that r , the number of terms in the sequence of subgraphs of $G_{n,d/n}$, is upper bounded by $E(G_{n,d/n})$. Thus, the above inequality implies that

$$Pr \left[r \geq (1 + n^{-1/3}) \frac{dn}{2} \right] \leq \exp \left(-n^{1/4} \right).$$

The lemma follows. ◇

C.7 Proof of Corollary 1

The existence of the bijection T implies that $|\Omega_1| = |\Omega_2|$. Thus $\forall \xi \in \Omega_1$ it holds that

$$Pr[X = \xi] = Pr[T(X) = T(\xi)] = \frac{1}{|\Omega_1|}.$$

Since, for every $\sigma \in \Omega_2$ there is a unique $\sigma' \in \Omega_1$ such that $T(\sigma') = \sigma$ we get that

$$Pr[T(X) = \sigma] = \frac{1}{|\Omega_1|} = \frac{1}{|\Omega_2|}.$$

The corollary follows.